A Revisit of Hashing Algorithms for Approximate Nearest Neighbor Search

Deng Cai
The State Key Lab of CAD&CG
College of Computer Science, Zhejiang University, China
dengcai@cad.zju.edu.cn

arXiv:1612.07545v2 [cs.CV] 27 Dec 2016

ABSTRACT

Approximate Nearest Neighbor (ANN) search is a fundamental problem in many areas of machine learning and data mining. During the past decade, numerous hashing algorithms are proposed to solve this problem. Every proposed algorithm claims outperforms other state-of-the-art methods. However, there are serious drawbacks in the evaluation of existing hashing papers and most of the claims in these papers should be re-examined. 1) Most of the existing papers failed to correctly measure the search time which is essential for the ANN search problem. 2) As a result, most of the papers report the performance increases as the code length increases, which is wrong if we measure the search time correctly. 3) The performance of some hashing algorithms (*e.g.*, LSH) can easily be boosted if one uses multiple hash tables, which is an important factor should be considered in the evaluation while most of the papers failed to do so. In this paper, we carefully revisit many popular hashing algorithms and suggest one possible promising direction. For the sake of reproducibility, all the codes used in the paper are released on Github, which can be used as a testing platform to fairly compare various hashing algorithms.

1. INTRODUCTION

Nearest neighbor search plays an important role in many applications of machine learning and data mining. Given a dataset with N entries, the cost for finding the exact nearest neighbor is $O(N)$, which is very time consuming when the data set is large. So people turn to Approximate Nearest Neighbor (ANN) search in practice [1, 22]. Hierarchical structure (tree) based methods, such as KD-tree [3] ,Randomized KD-tree [35], K-means tree [7], are very popular methods to solve the ANN search problem. These methods perform very well when the dimension of the data is relatively low. However, the performance decreases dramatically as the dimension of the data increases [35].

During the past decade, hashing based ANN search methods [9, 39, 25] received considerable attentions. These methods generates binary codes for high dimensional data points (real vectors) while try to preserve the similarity among the original real vectors. A hashing algorithm generating one bit code actually partitions the original feature space into two parts, the points in one part receive code 1 and the points in the other part receive code 0. When l-bits code is used, the hashing algorithm actually partitions the feature space into 2^l parts, which can be named as hash buckets. Thus, all the data points fall into different hash buckets (associated with different binary codes). Ideally, if neighbor vectors fall into the same bucket or the nearby buckets (measured by the hamming distance of binary codes), the hashing based methods can efficiently retrieve the nearest neighbors of a query point.

Different hashing algorithms partition the feature space based on different criterion. One of the most popular hashing algorithm is Locality Sensitive Hashing (LSH) [15, 9]. LSH is fundamentally based on random projection, *i.e.*, it uses random projection to partition the feature space. Some other popular hashing algorithms include Spectual Hashing [39], Kernelized LSH [25], Iterative Quantization [10]. Please see [37] for a detailed survey on various hashing algorithms.

Every new proposed hashing algorithm claims outperforms other state-of-the-art algorithms in their papers. However, there are serious drawbacks in the evaluation of existing hashing papers:

- All the ANNS methods sacrifice accuracy for speed. However, most of the existing hashing papers [39, 25, 24, 36, 42, 12, 10, 20, 41, 30, 34, 16, 14, 38, 23, 26, 13, 8, 17, 40, 21, 43, 28] only report some accuracy measures (*i.e.*, mean average precision, precision at K samples, precision-recall curves *etc.*). They failed to report what is the search time to achieve such accuracy.

- As a result, many hashing papers [39, 36, 10, 20, 30, 34, 16, 14, 38, 23, 40, 8, 26, 21, 43, 28] report the performance of a hashing algorithm increases as the code length increases, which is wrong if we measure the search time correctly.

- The performance of some hashing algorithms (*e.g.*, LSH [9]) can easily be boosted if one uses multiple hash tables, while it is not clear how to use the same trick (multiple tables) for some other algorithms (*e.g.*, Spectral Hashing [39]). This is an important factor should be considered in the evaluation while most of the papers failed to do so.

Thus, the claims in most of the existing hashing papers should be re-examined.

In this paper, we verified the drawbacks in the evaluations of most existing hashing papers by carefully designed experiments. Then we carefully revisit many popular hashing methods and suggest one possible promising direction. For the sake of reproducibility, all the codes used in the paper are released on Github[1][2]. This

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Submission to KDD'17 August 13 - 17, 2017, Halifax, Nova Scotia, Canada

© 2016 ACM. ISBN 978-1-4503-2138-9.

DOI: 10.1145/1235

[1]https://github.com/dengcai78/MatlabFunc/tree/master/ANNS/Hashing
[2]https://github.com/fc731097343/efanna/tree/master/samples_hashing

Algorithm 1 Search with Hash Index

Input: base set D, query vector q, hash index (binary codes for all the points in the base set), the pool size P and the number K of required nearest neighbors,

Output: K points from the base set which are close to q.

1: Encode the query q to the binary code b;
2: Let hamming radius $r = 0$;
3: Let candidate pool set $C = \emptyset$;
4: **while** $|C| < P$ **do**
5: Get the buckets list L with hamming distance r to b;
6: **for all** bucket bb in L **do**
7: Put all the points in bb into C;
8: **if** $|C| >= P$ **then**
9: Keep the first P points in C;
10: **goto** 15;
11: **end if**
12: **end for**
13: $r = r + 1$;
14: **end while**
15: return the closet K points to q in C.

can be used as a testing platform to fairly compare various hashing methods.

2. SEARCH WITH HASHING INDEX

All the hashing methods transform the real vectors to binary codes. The binary codes then can be used as a hashing index for on-line search. The general procedure of searching with hashing index is as follows (Suppose the user submit a query q and ask for K neighbors of the query):

1. The search system encode the query to the binary code b.
2. The search system finds P binary codes in the index which are closest to b based on the hamming distance.
3. The search system performs a linear scan within the P points, return the top K points which are closet to q to the user.

We summarized this procedure in Algorithm 1.

Based on this search procedure, we can see that the time on searching with a hashing index contains three parts [19]:

1. **Coding** time: the line 1 in Algorithm 1.
2. **Locating** time: the lines 2 $\sim$ 14 in Algorithm 1.
3. Linear **Scanning** time: the line 15 in Algorithm 1.

There is only one parameter P in this algorithm and one can tune this parameter to control the accuracy-efficiency trade-off.

3. WHAT'S WRONG WITH THE EVALUATION IN EXISTING HASHING PAPERS

After getting K neighbors of the query, we can use various accuracy measures to measure the performance of different hashing algorithms. Some well known measures include mean average precision (MAP), precision at K samples, precision at hamming radius r, recall at K samples, and precision-recall curves.

Most of the existing hashing papers [39, 25, 24, 36, 42, 12, 10, 20, 41, 30, 34, 16, 14, 38, 23, 26, 13, 8, 17, 40, 21, 43, 28] used these accuracy measures to evaluate the performance of various hashing algorithms. However, these accuracy measures are not enough. Remember, all the ANNS methods sacrifices accuracy for efficiency. If one does not care about the search time, then brute-force search can achieve the best performance. One can design a very simple hashing algorithm: encoding all the real vectors to exactly the same binary code. If this simple hashing algorithm is used, searching with the hashing index becomes exactly same as the brute-force search. However, if we only use the above mentioned accuracy measures for evaluation. This simple hashing algorithm will ranked No.1 among all the hashing algorithms.

3.1 Correctly Evaluate The Search Time

Thus, the search time must be reported together with the accuracy measures. The **accuracy-time** curve is a good choice to evaluate the performance of various hashing algorithms.

Some hashing papers do have reported the search time in some extent:

- [30, 29, 28, 18] reported the training and testing time of various hashing algorithms. In an ANNS system, the training time of a hashing algorithm is just the indexing time, while the testing time of a hashing algorithm is simply the **coding** time we mentioned in the previous section. For almost all the popular hashing algorithms, this coding time can be neglected comparing with locating time and linear scanning time. Please see Table 3 and Figure 7 for details.
- [25, 36, 12, 10, 41, 30, 34, 14, 29, 13, 8, 17, 40, 21, 43] reported **accuracy-# of retrieved samples** curves of various hashing algorithms. **# of retrieved samples** does related to the search time. However, it only reflects the linear **scanning** time we mentioned in the previous section, *i.e.*, the **scanning** time will linearly increases as the # of retrieved samples increases. Actually, the # of retrieved samples is exactly the parameter P in Algorithm 1.

In most of the existing hashing papers, the **locating** time is ignored, which is the main drawback in the evaluation of all these papers.

Given a binary code b, locating the hashing bucket corresponding to b costs $O(1)$ time (by using std::vector or std::unordered_map). If we only need to locate a small number of hash buckets, the locating time can be neglected. This happened if neighbor vectors fall into the same bucket or the nearby buckets (measured by the hamming distance of binary codes) ideally. However, there is no guarantee that all the neighbor vectors will fall into the nearby buckets. To ensure a high *recall* (the number of true neighbors within the returned points set divides by the number of required neighbors), one needs to examine many hashing buckets (*i.e.*, enlarge the search radius in hamming space).

Given a l-bits binary code b, considering those hashing buckets whose hamming distance to b is small or equal to r. It is easy to show the number of these buckets is $\sum_{i=0}^{r} \binom{l}{i}$, which increases almost exponentially with respect to i. Thus, the locating time can not be ignored if we aim at achieving a high recall (*i.e.*, we need to examine many hash buckets).

3.1.1 Long Codes Means Better Performance?

Many existing hashing papers [39, 36, 10, 20, 30, 34, 16, 14, 38, 23, 8, 26, 40, 21, 43, 28] reported that the performance of a hashing method increases as the code length increases. This claim only holds when the search time is not considered or only the linear scanning time is considered (by using **accuracy-# of retrieved samples** curves). If we correctly measure the search time (consider both locating time and scanning time), we will get a completely different conclusion.

To verify this, we designed the experiments as follows:

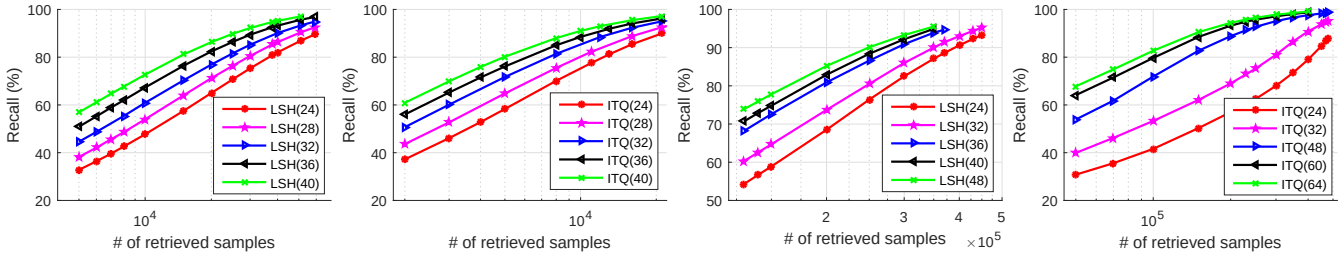

(a) LSH on SIFT1M (b) ITQ on SIFT1M (c) LSH on GIST1M (d) ITQ on GIST1M

Figure 1: The ANN search results of LSH and ITQ on two datasets. The y-axis denotes the average recall while the x-axis denotes the # of retrieved samples. We can clearly see longer codes means better performance from these figures.

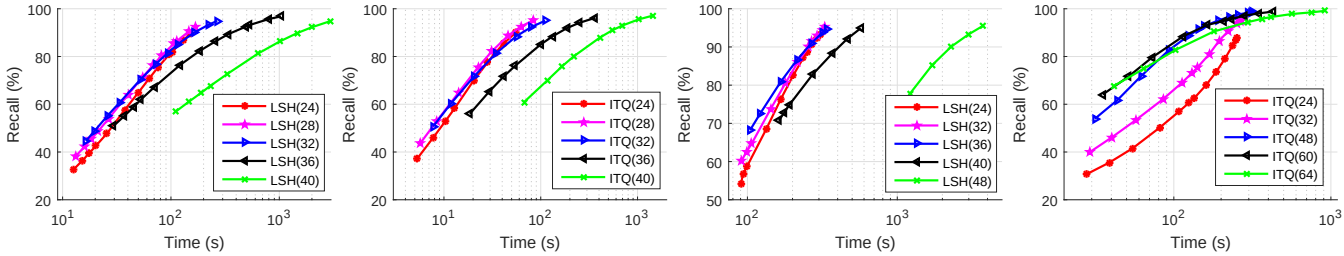

(a) LSH on SIFT1M (b) ITQ on SIFT1M (c) LSH on GIST1M (d) ITQ on GIST1M

Figure 2: The ANN search results of LSH and ITQ on two datasets. The y-axis denotes the average recall while the x-axis denotes the time. A hashing algorithm has an optimal code length which is different for different algorithms on different datasets.

Table 1: Data sets

data set	dimension	base number	query number
SIFT1M	128	1,000,000	10,000
GIST1M	960	1,000,000	1,000

- Two popular hashing algorithms are tested, they are Locality Sensitive Hashing (**LSH** for short) [9] and Iterative Quantization (**ITQ** for short) [10].

- Two popular datasets, SIFT1M and GIST1M, are used[3]. The basic statistics of these two datasets are shown in Table 1.

Given a query point, the algorithm is expected to return K points. Then we need to examine how many points in this returned set are among the true K nearest neighbors of the query. Suppose the returned set of K points given a query is R and the true K nearest neighbors set of the query is R', the $precision$ and $recall$ can be defined [31] as

$$precision(R) = \frac{|R' \cap R|}{|R|}, \quad recall(R) = \frac{|R' \cap R|}{|R'|}. \quad (1)$$

Since the sizes of R' and R are the same, the recall and the precision of R are actually the same. We fixed $K = 100$ throughout our experiments.

Figure 1 shows the **recall-# of retrieved samples** curves of LSH and ITQ on both datasets. Based on these figures, we can get the conclusion that longer codes means better performance. It means that long codes can generate better candidate points. This result is consistent with many existing hashing papers and easy to understand. Each bit is a partition of the feature space and same code (0 or 1) on this bit means two points are at the same side of this partition. Neighbors in hamming space with long code are more likely to be the true neighbors in the original space.

However, this is not the full story. Figure 2 shows the **recall-time**[4] curves of LSH and ITQ on both datasets. We can clearly see that the search time significantly increases as the code length increases. Although longer code can generate better candidates, they need more time to locate these candidates. Thus, it becomes practically useless for very long codes. A hashing algorithm has an optimal code length which is different for different algorithms on different datasets. For a million-scale dataset, the optimal code length is usually smaller than 64.

From Figure 1 (a), we can see if # of retrieved samples is set as 5,000 (the starting point of all the curves), LSH (24 bits) reaches around 30% recall while LSH (40 bits) reaches almost 60% recall. From Figure 2 (a), we can see LSH (24 bits) uses around 10s while LSH (40 bits) uses around 100s. Since # of retrieved samples is set as 5,000, the linear scanning times for the two cases are the same. It is the locating time causing this 90s difference, *i.e.*, the locating time of LSH (40 bits) is about 10 times than that of LSH (24 bits).

Figure 3 provides an explanation. The first row of Figure 3 shows the number of queries (the total number is 10,000) which successfully located 5,000 samples in different hamming radius of LSH with different code length. Figure 3 (a) shows more than 5,000 queries located 5,000 samples successfully within hamming radius 0 (*i.e.*, more than 5,000 queries only need to visit one hash bucket) when LSH uses 24 bits. As a comparison, if LSH uses 40 bits, Figure 3 (e) shows more than 4,000 queries need to expand the

[3]Both of two datasets can be downloaded at http://corpus-texmex.irisa.fr/

[4]We only record the locating time and scanning time here. The coding time can be neglected comparing to locating time and scanning time. Please see Table 3 and Figure 7 for details.

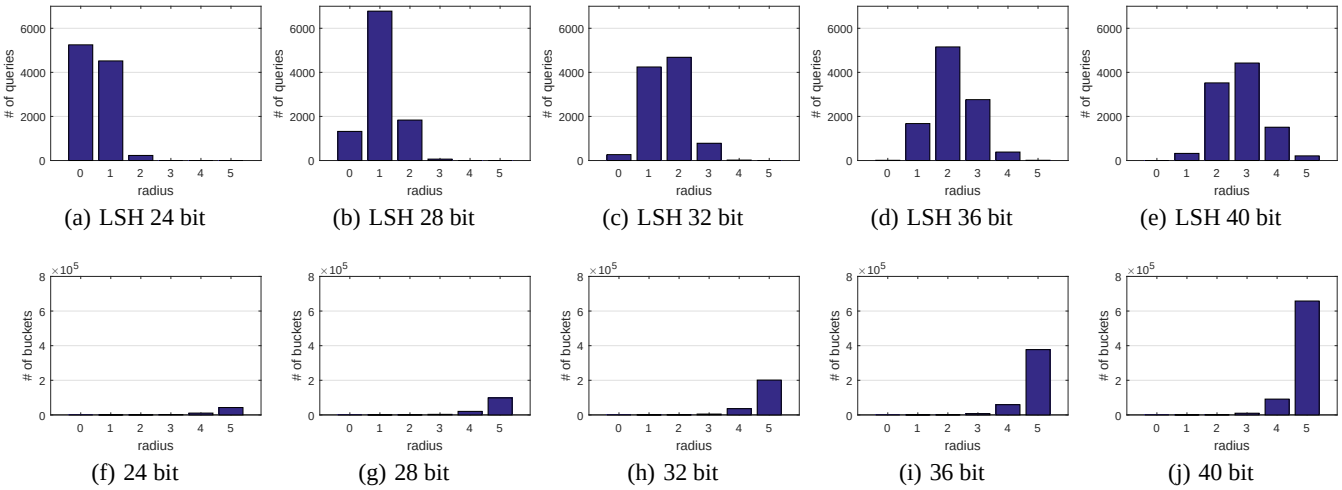

(a) LSH 24 bit (b) LSH 28 bit (c) LSH 32 bit (d) LSH 36 bit (e) LSH 40 bit

(f) 24 bit (g) 28 bit (h) 32 bit (i) 36 bit (j) 40 bit

Figure 3: The first row shows the number of queries (the total number is 10,000) which successfully located 5,000 samples in different hamming radius of LSH with different code length on SIFT1M. The second row shows that the hashing buckets number grows quickly as the code length increases when the radius r fixed. This figure explains why the search time grows quickly as the code length increases.

hamming radius to 3 to locate 5,000 samples (*i.e.*, more than 4,000 queries need to visit around $\sum_{i=0}^{3}\binom{40}{i} = 10,701$ hash buckets). The second row of Figure 3 shows that the hash buckets number grows quickly as the code length increases when the radius r fixed. This figure explains why the search time grows so quickly as the code length increases. For a million-scale data set, code length longer than 64 bits has no practical meaning.

3.2 Evaluation with Multiple Hash Tables

Our analysis in the previous section shows that the locating time should not be neglected. Actually, when aiming at a high recall (need to visit many hash buckets), the locating time is dominant on the search time. One simply way to reduce the locating time (while keep the high recall level) is using multiple hash tables.

Take LSH as an example. Since LSH is essentially based on random projection, two hash tables generate by LSH naturally will be different (*i.e.*, a query point will have different neighbor vectors in nearby hash buckets). To locate P points, if we only have one hash table, we have to increase the hamming radius r if the points in all the buckets within the hamming distance r are not enough. This will increase the locating time a lot. If we have multiple hash tables, instead of increasing r, we can scan the buckets within the hamming radius r in all the tables, which gives us a larger chance to locate enough data points.

Figure 4 shows both the **recall-# of retrieved samples** and **recall-time** curves of LSH and ITQ with different number of tables on the two datasets. The results are very consistent: both two algorithms gain advantage by using multiple tables on both datasets. Using multiple tables can generate better candidate points, just as the long codes does. Moreover, using multiple tables does not increase the locating time. Multiple tables is a better strategy overwhelming long codes. Actually, we used 16 tables in all the previous experiments. The performance improvement from using 1 table to 4 tables is biggest. When we increase the number of tables further, the performance gradually increased.

Figure 5 shows the number of queries (the total number is 10,000) which successfully located 5,000 samples in different hamming radius of LSH with different number of tables on SIFT1M. We can see that with more tables, queries are more likely to locate enough neighbor points within small hamming radius. This is the reason why using multiple hash tables does not increase (in many case actually decrease) the locating time (Although one needs to scan multiple tables, the time spent on each table becomes small).

To use the multiple tables trick, a hashing algorithm must generate different hash tables for the same dataset. Some hashing algorithms (*e.g.*, LSH and ITQ) have randomness in nature and naturally can use the multiple tables trick. Some other hashing algorithms (*e.g.* Anchor Graph Hashing [30] and Complementary Projection Hashing [17]) are essentially deterministic. However, these algorithms use the idea of anchors (landmarks) and different anchors will result different hash tables. Thus, these algorithms can also use the multiple tables trick.

There are still many hashing algorithms [42, 36, 34, 43] can not use this multiple tables trick (*i.e.*, these algorithms will generate the same hash table for the same dataset every time). Thus it is not fair to compare all these hashing algorithms using only one table.

[19] reported the inferior performance of hashing algorithms comparing to flann's KD-tree [33]. However, [19] used randomized KD-trees with 16 trees but only used one table for hashing algorithms. Thus, the comparisons in [19] is not fair to hashing algorithms and the conclusion should be re-examined.

4. A THROUGH COMPARISON

Our analysis in the previous section shows that the evaluations in most of the existing hashing papers have serous drawbacks. In this section, we aim at providing a through comparison of 12 popular hashing algorithms on SIFT1M and GIST1M.

4.1 Compared Algorithms

Twelve popular hashing algorithms compared in the experiments are listed as follows:

- **LSH** is a short name for Locality Sensitive Hashing [9]. LSH is based on random projection and is frequently used as a baseline method in various hashing papers.
- **SH** is a short name for Spectral Hashing [39]. SH is based on quantizing the values of analytical eigenfunctions computed along PCA directions of the data.

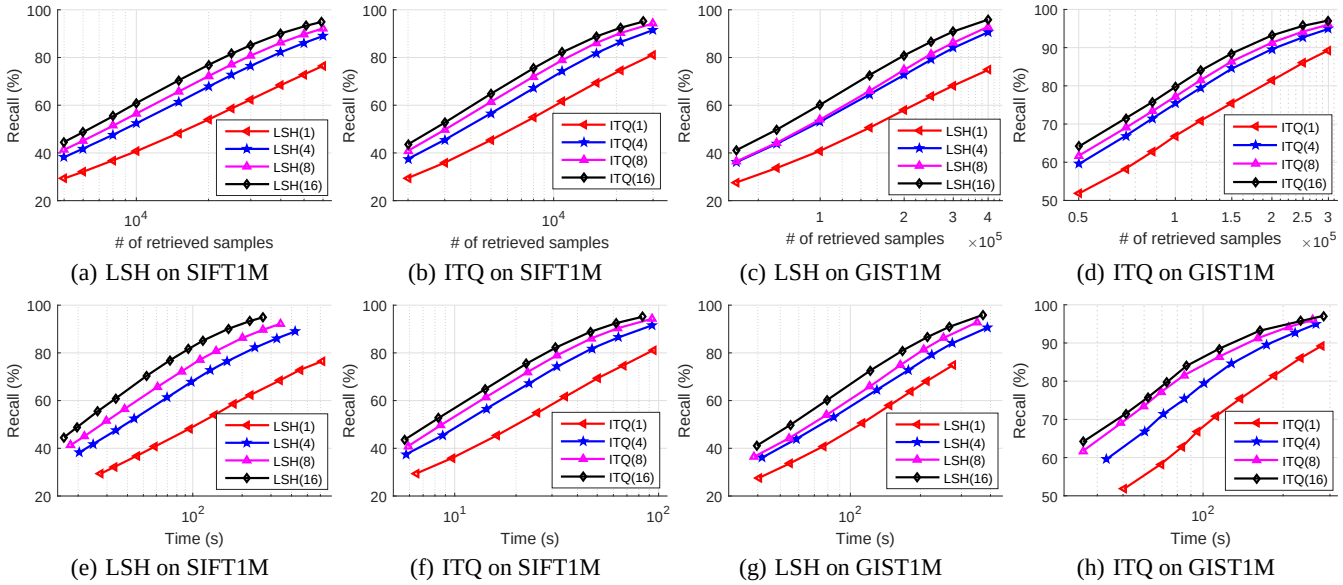

(a) LSH on SIFT1M (b) ITQ on SIFT1M (c) LSH on GIST1M (d) ITQ on GIST1M

(e) LSH on SIFT1M (f) ITQ on SIFT1M (g) LSH on GIST1M (h) ITQ on GIST1M

Figure 4: The first row shows the recall-# of retrieved samples curves of LSH and ITQ on both two datasets and the second row shows the recall-time curves. We can see that multiple tables not only generates better candidate neighbors but also requires short locating time.

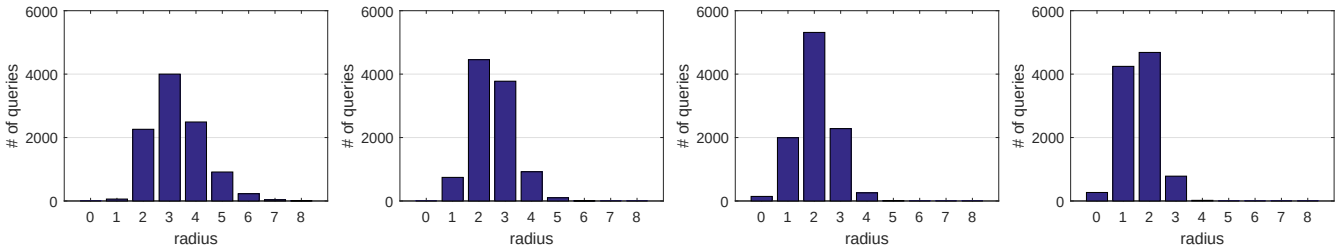

(a) 1 table (b) 4 table (c) 8 table (d) 16 table

Figure 5: The number of queries which successfully located 5,000 samples in different hamming radius of LSH with different number of tables on SIFT1M (the total query number is 10,000).

- **KLSH** is a short name for Kernelized Locality Sensitive Hashing [25]. KLSH generalizes the LSH method to the kernel space.
- **BRE** is a short name for Binary Reconstructive Embeddings [24].
- **USPLH** is a short name for Unsupervised Sequential Projection Learning Hashing [36].
- **ITQ** is a short name for ITerative Quantization [10]. ITQ finds a rotation of zero-centered data so as to minimize the quantization error of mapping this data to the vertices of a zero-centered binary hypercube.
- **AGH** is a short name for Anchor Graph Hashing [30]. It aims at performing spectral analysis [2] of the data which shares the same goal with Self-taught Hashing [42]. The advantage of AGH over Self-taught Hashing is the computational efficiency. AGH uses an anchor graph [27] to speed up the spectral analysis. There are two versions of AGH: one layer AGH and two layer AGH. We use **AGH1** and **AGH2** to denote them.
- **SpH** is a short name for Spherical Hashing [14]. SpH uses a hyperspherebased hash function to map data points into binary codes.
- **IsoH** is a short name for Isotropic Hashing [23]. IsoH learns the projection functions with isotropic variances for PCA projected data. The main motivation of IsoH is that PCA directions with different variance should not be equally treated (one bit for one direction).
- **CH** is a short name for Compressed Hashing [26]. CH first learns a landmark (anchor) based sparse representation [5] then followed by LSH.
- **CPH** is a short name for Complementary Projection Hashing [17]. CPH finds the projections sequentially by balancing crossing the sparse region and the points evenly distributed in each hash bucket.
- **DSH** is a short name for Density Sensitive Hashing [18]. DSH finds the projections which aware the density distribution of the data.

The multiple tables trick can not be used for SH and USPLH, we

Table 2: Indexing Time (s)

Method	SIFT1M	GIST1M	Method	SIFT1M	GIST1M
LSH	3.8	15.4	SH*	7.1	15.5
KLSH	795.0	2080.9	BRE	405.7	344.0
USPLH*	33.5	337.3	ITQ	358.8	428.7
AGH1	1504.2	2883.1	AGH2	1645.5	3006.6
SpH	824.7	1979.0	IsoH	11.2	131.7
CH	1163.8	2502.5	CPH	15876.0	152552.8
DSH	68.1	285.9			

*These two algorithms learn one table while the others learn 16 tables.

Table 3: Coding Time (s)

Method	SIFT1M	GIST1M	Method	SIFT1M	GIST1M
LSH	0.032	0.016	SH*	0.118	0.020
KLSH	2.086	0.465	BRE	0.539	0.288
USPLH*	0.002	0.001	ITQ	0.032	0.016
AGH1	4.652	0.964	AGH2	5.016	1.008
SpH	0.100	0.049	IsoH	0.032	0.016
CH	4.747	0.908	CPH	2.963	0.567
DSH	0.032	0.016			

*These two algorithms use one table while the others use 16 tables.

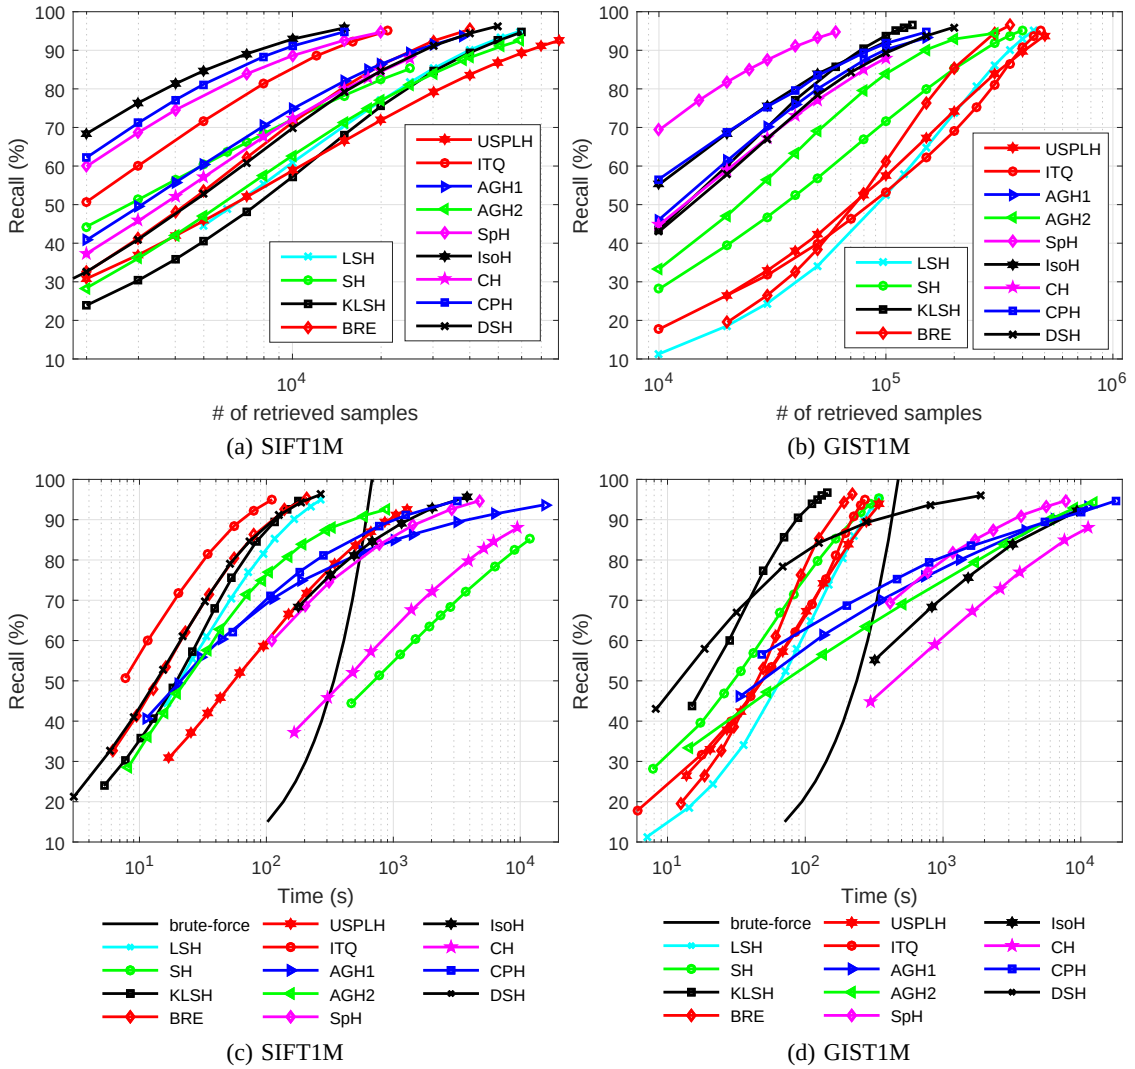

(a) SIFT1M (b) GIST1M

(c) SIFT1M (d) GIST1M

Figure 6: The ANN search results of various hashing algorithms with 32 bits codes on SIFT1M and GIST1M. (a) (b) are recall-# of retrieved samples curves and (c) (d) are recall-time curves.

simply use one hash table for these two algorithms. For the other algorithms, we use 16 tables throughout the experiments.

For KLSH, AGH, CH and CPH, we need to pick a certain number of anchor (landmark) points. We use the same 1,000 anchor points for all these algorithms and the anchor points are generated by using kmeans with 5 iterations. Also, the number of nearest anchors is set to 50 for AGH and CH algorithms. Please see [4] for details.

All the hashing algorithms are implemented in Matlab and we use these hashing functions to learn the binary codes for both base

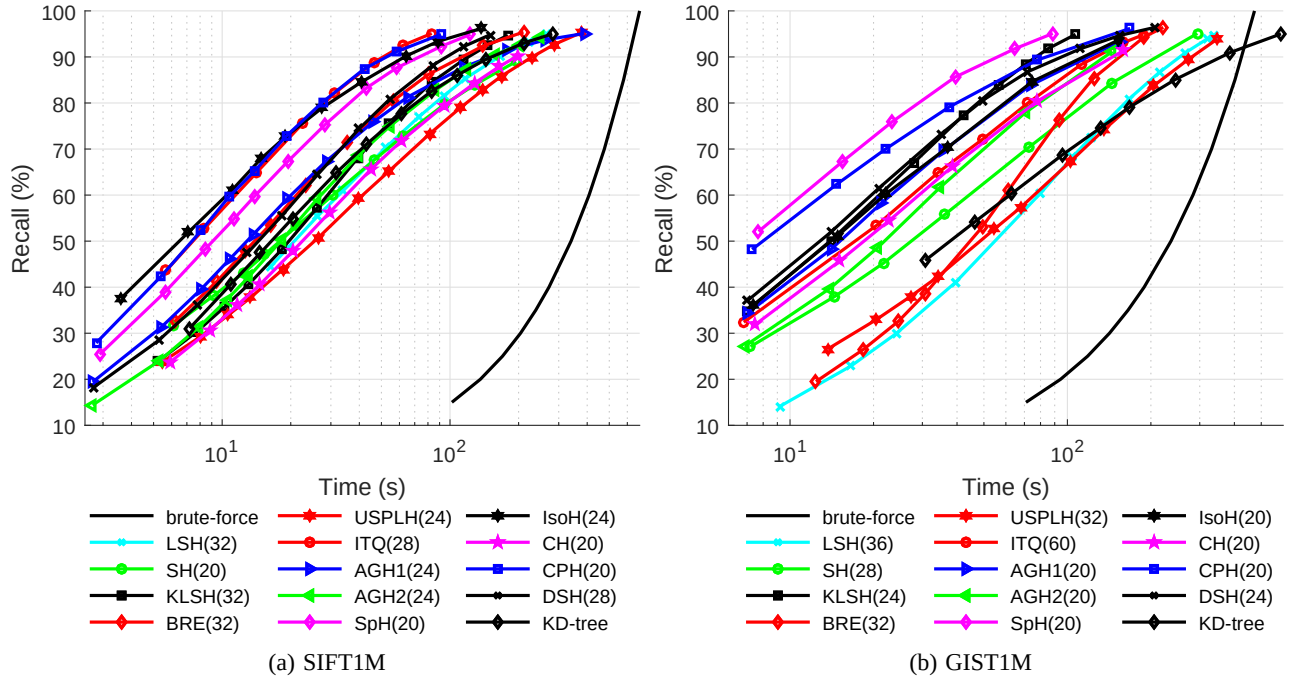

(a) SIFT1M (b) GIST1M

Figure 7: The recall-time curves of various hashing algorithms with the optimal code length on SIFT1M and GIST1M. The number in parenthesis after the algorithm name is the code length.

vectors and query vectors. The binary codes can then be fed into Algortihm 1 for search. The Matlab functions are run on a i7-5930K cpu with 128G memory and the Algortihm 1 is implemented in c++ and run on a i7-3770K CPU and 32G memory. For the sake of reproducibility, both the Matlab codes and c++ codes are released on Github.

Besides all the hashing algorithms, we also report the performance of brute-force search and flann's KD-tree.

- **brute-force**. The performance of brute-force search is reported to show the advantages of using hashing methods. To get different recall, we simply perform brute-force search on different percentage of the query number. For example, the brute-force search time on 90% queries of the origin query set stands for the brute-force search time of 90% recall.

- **KD-tree**. Flann [33] is a well-known open source library for ANN search. It integrates Randomized KD-tree [35] and Kmeans tree [7]. In our experiment, we use the randomized KD-tree and also set the trees number as 16.

4.2 Results on 32 Bits Code

We first try all the hashing algorithms with 32 bits code. Since we use Matlab functions to learn the binary codes and c++ algorithm to search with the hashing index, the coding time and later time (locating time and scanning time) can not be added together and we reported them separately.

Table 2 and 3 report the indexing (learning) time and coding (testing) time of various hashing algorithms. The indexing (learning) stage is performed off-line thus the indexing time is not very crucial as long as the indexing (hashing learning) time is not too long. CPH might be the only algorithm which is not practical because the indexing time is extremely long. The coding time of all the hashing algorithms are short. Even for the slowest algorithm, AGH2, the coding time is only 5s on SIFT1M and 1s on GIST1M. Thus the coding time can be ignored compared with locating time and scanning time (see Figure 6) when we consider the total search time. The coding times of KLSH, AGH, CH and CPH are significantly longer than other algorithms simply because these algorithms need to compute the distance of a query point to the anchor points (1000 anchor points).

Figure 6 shows the **recall-# of retrieved samples** and **recall-time** curves of various hashing algorithms. Based on the **recall-# of retrieved samples** curves, we can get the similar conclusions as many previous hashing papers [17, 18]. The results in Figure 6 (a)(b) are better than the results reported in [17] simply because we used 16 hash tables while [17] only use one hash table. However, **recall-# of retrieved samples** curve are misleading that the algorithms performing good on **recall-# of retrieved samples** curve actually perform really bad (many algorithms are even worse than the brute-force search) on **recall-time** curve. The latter actually has more practical meaning.

However, we still can not get the conclusion that which hashing algorithm performs best simply based on **recall-time** curves in Figure 6. There are reasons to doubt that 32 bits might not the optimal code length for those algorithms which perform really good on **recall-# of retrieved samples** curve but perform bad on **recall-time** curve. It is very natural that different hashing algorithm has different optimal code length on different dataset. It is unfair to compare different algorithms under the same code length, just as many of the existing hashing papers do.

4.3 Results on The Optimal Code Length

We now try all the hashing algorithms with their optimal code length. The recall-tTheime curves of all the hashing algorithms together with that of flann's KD-tree and brute-force are shown in the Figure 7. We can clearly see that all the hashing algorithms gain some advantages over the brute-force search and many hashing algorithms are better than flann's KD-tree. This result is contradict

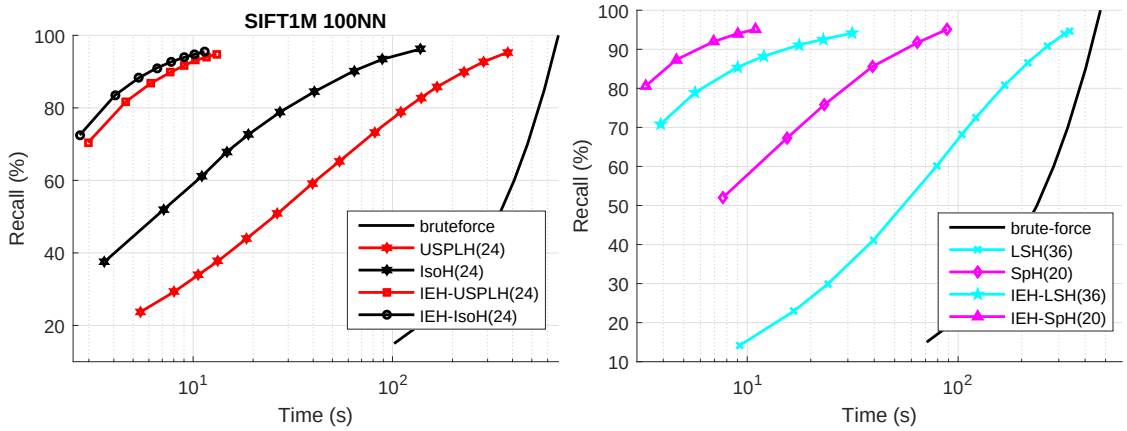

(a) SIFT1M (b) GIST1M

Figure 8: The ANN search results of IEH extension and its corresponding hashing algorithms. The number in parenthesis after the algorithm name is the code length of the hashing algorithm. The IEH extension boosts its corresponding hashing algorithm significantly, no matter which hashing algorithm is used.

with the previous paper [19] simply because we used 16 tables for hashing algorithms and 16 trees for KD-tree while [19] used one table for hashing algorithm and 16 trees for KD-tree.

On SIFT1M, the top 4 algorithms are IsoH, CPH, ITQ and SpH. These four algorithms are significantly better than other algorithms. On GIST1M, the best algorithm is SpH, followed by CPH, then KLSH and DSH. Generally speaking, the optimal code lengths of different algorithms are different on different datasets. However, noticing that the optimal code length of many algorithms (SH, SpH, CH, CPH on SIFT1M and AGH, SpH, IsoH, CH, CPH on GIST1M) is 20, and there are $2^{20} = 1,048,576$ buckets in total if the code length is 20. The buckets number is approximately equal to the data size. This might be a reasonable way to estimate the optimal code length. But it needs further exploration.

Considering the performance on both data sets, SpH is the recommended algorithm. Despite its good performance, We do not recommend CPH simply for its very long indexing (training) time.

5. ITERATIVE EXPANDING HASHING

Given the Figure 7, we are still not satisfied with the performance of hashing algorithms. If a high recall (say, 95%) is needed, even the best hashing algorithm can only obtain 6-times speed up over the brute-force search. How can we boost the performance of an existing hashing algorithm?

Recently, graph based techniques have drawn considerable attention [11, 19, 32, 6]. The main idea of these methods is *a neighbor of a neighbor is also likely to be a neighbor*. At offline stage, they need to build a kNN graph, which can be regarded as a big table recording the top k closest neighbors of each point in database. At online stage, given a query point, they first assign the query some points as initial candidate neighbors, and then check the neighbors of the neighbors iteratively to locate closer neighbors. The Iterative Expanding Hashing (IEH) [19] is the most related work to hashing algorithms. IEH uses hashing algorithms to generate the initial candidate neighbors and Algorithm 2 gives the main routine of IEH.

To use IEH, one needs to build a kNN graph at the indexing stage. Since building an exact kNN graph is very time consuming, we use Efanna [6], an open source library to build an approximate kNN graph. More specifically, Efanna spent 98s for a 98% accuracy 50NN graph on SIFT1M and 404s for a 90% accuracy 50NN graph on GIST1M. These two graphs are used in our experiment.

Algorithm 2 IEH Search Algorithm

Input: base set D, query vector q, hash index, the kNN graph G, the pool size P, the expansion size E, the iteration number I and the number K of required nearest neighbors.
Output: K points from the base set which are close to q.

```
Search for E neighbors of q using algorithm 1, form the initial set C.
iter = 0
while iter < I do
    Let candidate set CC = ∅
    for all point n in C do
        S_n is the neighbors of point n based on G.
        for all point nn in S_n do
            Put nn into CC.
        end for
    end for
    Move all the points in CC to C;
    Keep E points in C which are closest to q.
    iter = iter + 1
end while
return the closet K points to q in C.
```

IEH has three additional essential parameters, the pool size P, the expansion size E, the iteration number I. We fixed $I = 6$ and set $P = 200$ on SIFT1M and $P = 2000$ on GIST1M. Then we change the expansion size to get the **recall-time** curves.

Figure 8 shows the ANN search results of IEH extension and its corresponding hashing algorithms. On both two datasets, we pick the best and worst hashing algorithm based on the Figure 7. The IEH extension boosts its corresponding hashing algorithm significantly (10-times speed up), no matter which hashing algorithm is used. Using IEH extension, the performance gap between different hashing algorithms becomes smaller. The conclusion is that the future direction of ANN search is using additional data structures (*e.g.*, kNN graph here).

5.1 The Best Hashing Algorithm for IEH

A natural question here is what is the best algorithm for IEH.

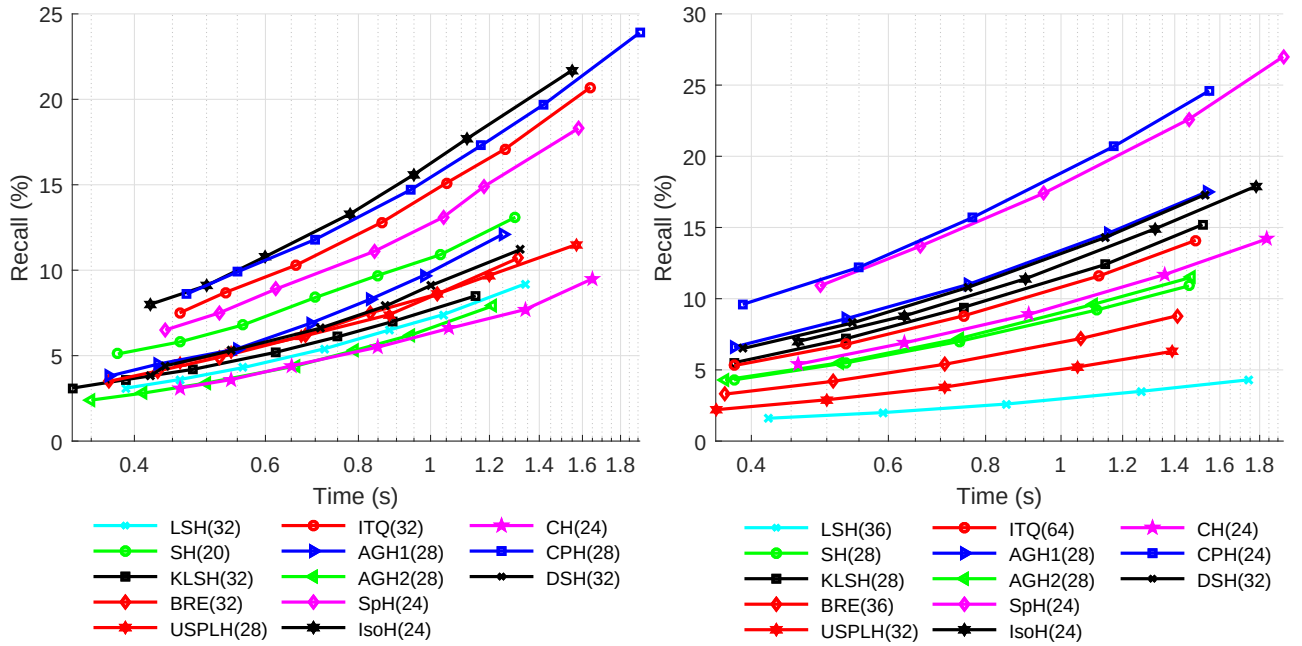

Figure 9: The recall-time curves of various hashing algorithms at the low recall region with the optimal code length on SIFT1M and GIST1M. The number in parenthesis after the algorithm name is the code length.

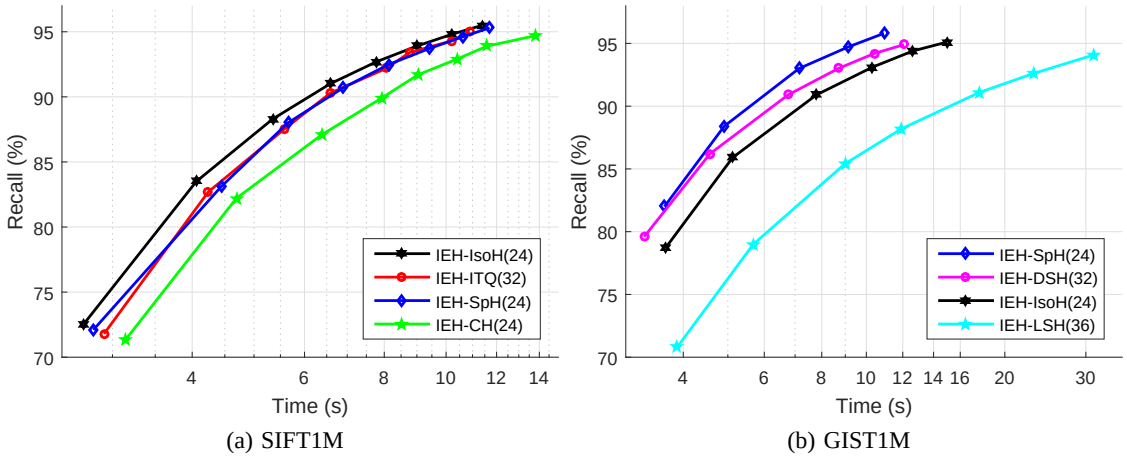

Figure 10: The ANN search results of IEH extension of various hashing algorithms. The number in parenthesis after the algorithm name is the code length of the hashing algorithm. The hashing algorithm has the best performance in the low recall region (see Figure 9) is best algorithm for IEH extension.

In Figure 8, we simply pick the hashing algorithms based on the Figure 7. However, IEH only needs a very small number of initial points (200 in SIGT1M and 2000 in GIST1M). Thus, the best hashing algorithm for IEH should produce very good initial points. In other words, we only care about the performance on the low recall region instead of the high recall region.

Figure 9 shows the recall-time curves of various hashing algorithms at the low recall region with the optimal code length on SIFT1M and GIST1M. Since the locating time plus the scanning time is around 1s, the coding time then can not be neglected. Thus, all the anchors based algorithms (KLSH, AGH, CH and CPH) should be removed from the best algorithm candidate list. On SIFT1M, the top three algorithms are IsoH, ITQ and SpH. On GIST1M, the top three algorithms are SpH, DSH and IsoH. The performance of SpH is significantly better than that of the other two. Overall, the recommended hashing algorithms for IEH are SpH and IsoH.

It is interesting to find that the optimal code length in Figure 9 is larger than the optimal code length in Figure 7. It is reasonable because on the low recall region one only needs to locate a small number of candidate points, making the locating time less important in the whole time.

Figure 10 shows the recall-time curves of IEH extensions of top three hashing algorithms and the worst algorithm on SIFT1M and GIST1M based on the Figure 9. The results agree with our initial judgment: The hashing algorithm has the best performance in the low recall region is best algorithm for IEH extension.

6. CONCLUSIONS

We carefully studied the evaluation of ANNS problem and verified the serious drawbacks in the evaluation of most of the exiting hashing papers. Then we revisit many popular hashing algorithms with extensive experiments and suggest one possible promising direction. All the codes used in the paper are released on Github, which can be used as a testing platform to fairly compare various hashing methods.

Acknowledgment

This work was supported in part by National Basic Research Program of China (973 Program) under Grant 2013CB336500 and National Youth Top-notch Talent Support Program. Any opinions, findings, and conclusions expressed here are those of the authors and do not necessarily reflect the views of the funding agencies.

7. REFERENCES

[1] S. Arya and D. M. Mount. Approximate nearest neighbor queries in fixed dimensions. In *Proceedings of the Fourth Annual ACM/SIGACT-SIAM Symposium on Discrete Algorithms*, 1993.

[2] M. Belkin and P. Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Advances in Neural Information Processing Systems 14*, 2001.

[3] J. L. Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, 1975.

[4] D. Cai and X. Chen. Large scale spectral clustering via landmark-based sparse representation. *IEEE Trans. Cybernetics*, 45(8):1669–1680, 2015.

[5] X. Chen and D. Cai. Large scale spectral clustering with landmark-based representation. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, 2011.

[6] C. Fu and D. Cai. Efanna : An extremely fast approximate nearest neighbor search algorithm based on knn graph. *arXiv:1609.07228*, 2016.

[7] K. Fukunaga and P. M. Narendra. A branch and bound algorithm for computing k-nearest neighbors. *IEEE Transactions on Computers*, 100(7):750–753, 1975.

[8] T. Ge, K. He, Q. Ke, and J. Sun. Optimized product quantization for approximate nearest neighbor search. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, 2013.

[9] A. Gionis, P. Indyk, and R. Motwani. Similarity search in high dimensions via hashing. In *Proceedings of the 25th International Conference on Very Large Data Bases*, 1999.

[10] Y. Gong and S. Lazebnik. Iterative quantization: A procrustean approach to learning binary codes. In *The 24th IEEE Conference on Computer Vision and Pattern Recognition*, 2011.

[11] K. Hajebi, Y. Abbasi-Yadkori, H. Shahbazi, and H. Zhang. Fast approximate nearest-neighbor search with k-nearest neighbor graph. In *Proceedings of the International Joint Conference on Artificial Intelligence*, volume 22, 2011.

[12] J. He, W. Liu, and S. Chang. Scalable similarity search with optimized kernel hashing. In *Proceedings of the 16th ACM International Conference on Knowledge Discovery and Data Mining*, 2010.

[13] K. He, F. Wen, and J. Sun. K-means hashing: An affinity-preserving quantization method for learning binary compact codes. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, 2013.

[14] J. Heo, Y. Lee, J. He, S. Chang, and S. Yoon. Spherical hashing. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, 2012.

[15] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing*, 1998.

[16] H. Jégou, M. Douze, and C. Schmid. Product quantization for nearest neighbor search. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(1):117–128, 2011.

[17] Z. Jin, Y. Hu, Y. Lin, D. Zhang, S. Lin, D. Cai, and X. Li. Complementary projection hashing. In *IEEE International Conference on Computer Vision*, 2013.

[18] Z. Jin, C. Li, Y. Lin, and D. Cai. Density sensitive hashing. *IEEE Trans. Cybernetics*, 44(8):1362–1371, 2014.

[19] Z. Jin, D. Zhang, Y. Hu, S. Lin, D. Cai, and X. He. Fast and accurate hashing via iterative nearest neighbors expansion. *IEEE transactions on cybernetics*, 44(11):2167–2177, 2014.

[20] A. Joly and O. Buisson. Random maximum margin hashing. In *The 24th IEEE Conference on Computer Vision and Pattern Recognition*, 2011.

[21] Y. Kalantidis and Y. S. Avrithis. Locally optimized product quantization for approximate nearest neighbor search. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, 2014.

[22] J. M. Kleinberg. Two algorithms for nearest-neighbor search in high dimensions. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing*, 1997.

[23] W. Kong and W. Li. Isotropic hashing. In *Advances in Neural Information Processing Systems 25*, 2012.

[24] B. Kulis and T. Darrell. Learning to hash with binary reconstructive embeddings. In *Advances in Neural Information Processing Systems 22, NIPS*, 2009.

[25] B. Kulis and K. Grauman. Kernelized locality-sensitive hashing for scalable image search. In *IEEE 12th International Conference on Computer Vision*, 2009.

[26] Y. Lin, R. Jin, D. Cai, S. Yan, and X. Li. Compressed hashing. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, 2013.

[27] W. Liu, J. He, and S.-F. Chang. Large graph construction for scalable semi-supervised learning. In *Proceedings of the 27th International Conference on Machine Learning*, 2010.

[28] W. Liu, C. Mu, S. Kumar, and S. Chang. Discrete graph hashing. In *Advances in Neural Information Processing Systems 27*, 2014.

[29] W. Liu, J. Wang, R. Ji, Y. Jiang, and S. Chang. Supervised hashing with kernels. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, 2012.

[30] W. Liu, J. Wang, S. Kumar, and S. Chang. Hashing with graphs. In *Proceedings of the 28th International Conference on Machine Learning*, 2011.

[31] J. Makhoul, F. Kubala, R. Schwartz, and R. Weischedel. Performance measures for information extraction. In *Proceedings of DARPA Broadcast News Workshop*, 2000.

[32] Y. A. Malkov and D. A. Yashunin. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *arXiv:1603.09320*, 2016.

[33] M. Muja and D. G. Lowe. Scalable nearest neighbor algorithms for high dimensional data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(11):2227–2240, 2014.

[34] M. Norouzi and D. J. Fleet. Minimal loss hashing for compact binary codes. In *Proceedings of the 28th International Conference on Machine Learning,ICML 2011*, 2011.

[35] C. Silpa-Anan and R. Hartley. Optimised kd-trees for fast image descriptor matching. In *Proceedings of the 2008 IEEE Conference on Computer Vision and Pattern Recognition*, 2008.

[36] J. Wang, S. Kumar, and S. Chang. Sequential projection learning for hashing with compact codes. In *Proceedings of the 27th International Conference on Machine Learning*, 2010.

[37] J. Wang, H. T. Shen, J. Song, and J. Ji. Hashing for similarity search: A survey. *arXiv preprint arXiv:1408.2927*, 2014.

[38] Y. Weiss, R. Fergus, and A. Torralba. Multidimensional spectral hashing. In *12th European Conference on Computer Vision*, 2012.

[39] Y. Weiss, A. Torralba, and R. Fergus. Spectral hashing. In *Advances in Neural Information Processing Systems 21, NIPS*, 2008.

[40] B. Xu, J. Bu, Y. Lin, C. Chen, X. He, and D. Cai. Harmonious hashing. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, 2013.

[41] H. Xu, J. Wang, Z. Li, G. Zeng, S. Li, and N. Yu. Complementary hashing for approximate nearest neighbor search. In *IEEE International Conference on Computer Vision*, 2011.

[42] D. Zhang, J. Wang, D. Cai, and J. Lu. Self-taught hashing for fast similarity search. In *Proceeding of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2010.

[43] K. Zhao, H. Lu, and J. Mei. Locality preserving hashing. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014.